Xen Project Summit

The journey towards using Xen in Embedded systems

Julien Grall

Whoami

- Part of the community since 2012
- Co-maintaining the Arm port
- Committer
- Currently employed by Amazon Web Services

Xen Architecture

- Type-1 hypervisor
- Domain controller (dom0)
- Full port for
  - X86
  - Arm-v7 A and Arm-v8 A
- On-ongoing port for
  - RISC-V
  - PPC

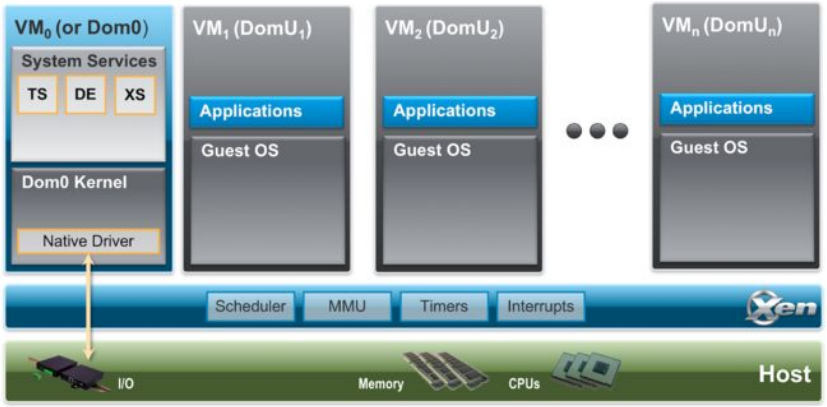

https://wiki.xenproject.org/wiki/Xen_Project_Software_Overview

Xen virtualization type

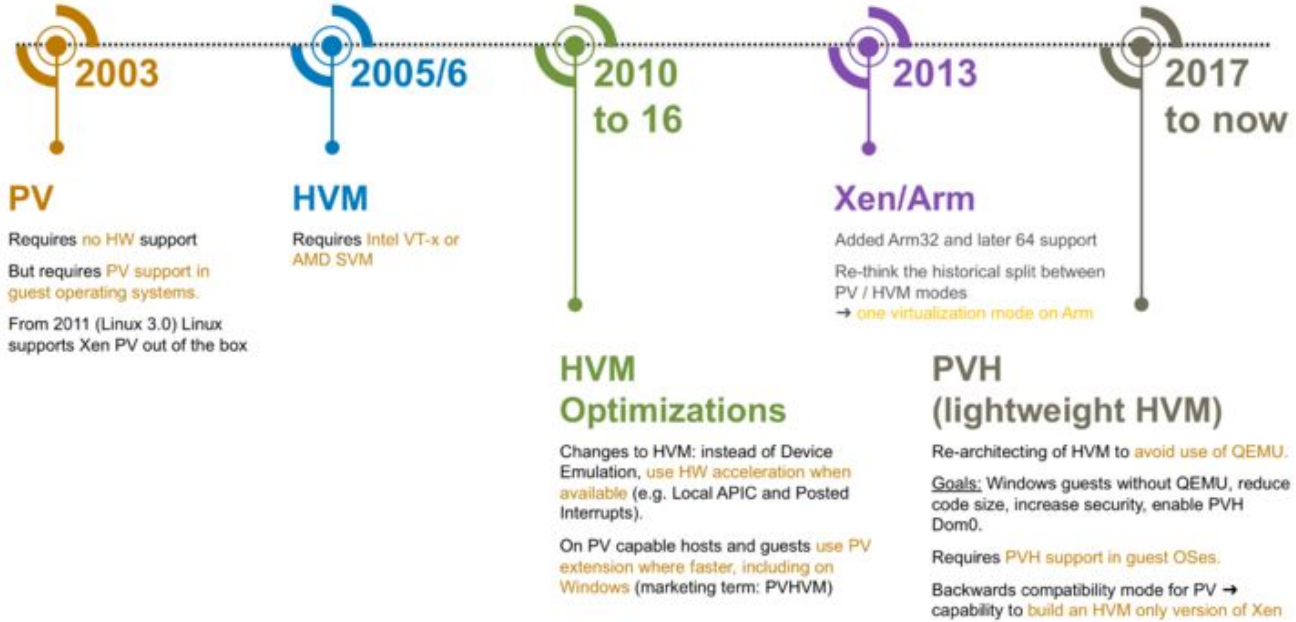

https://wiki.xenproject.org/wiki/Xen_Project_Software_Overview

Xen architecture (cont)

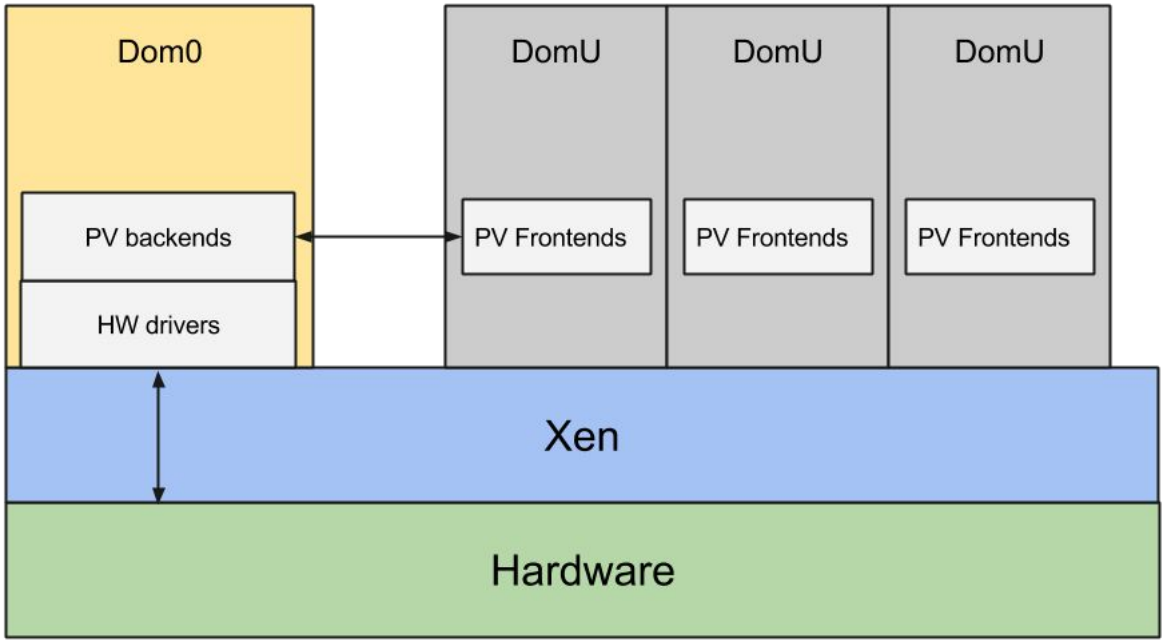

Driver domains

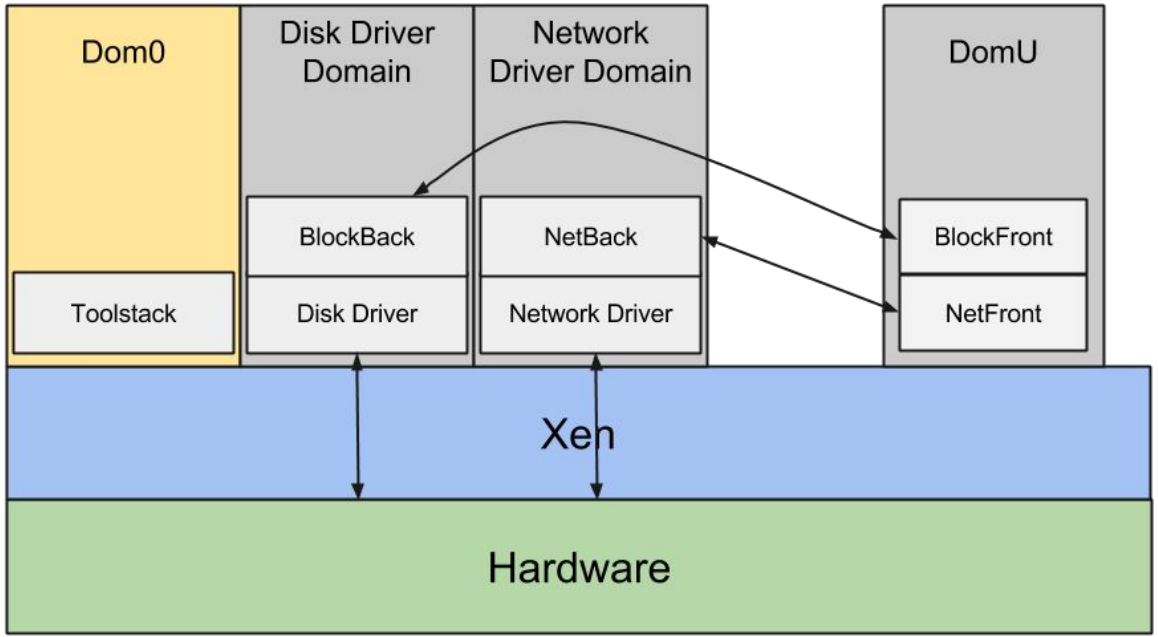

Scheduler

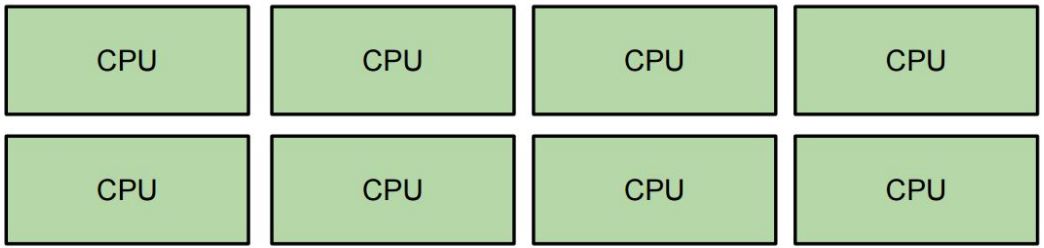

Scheduler

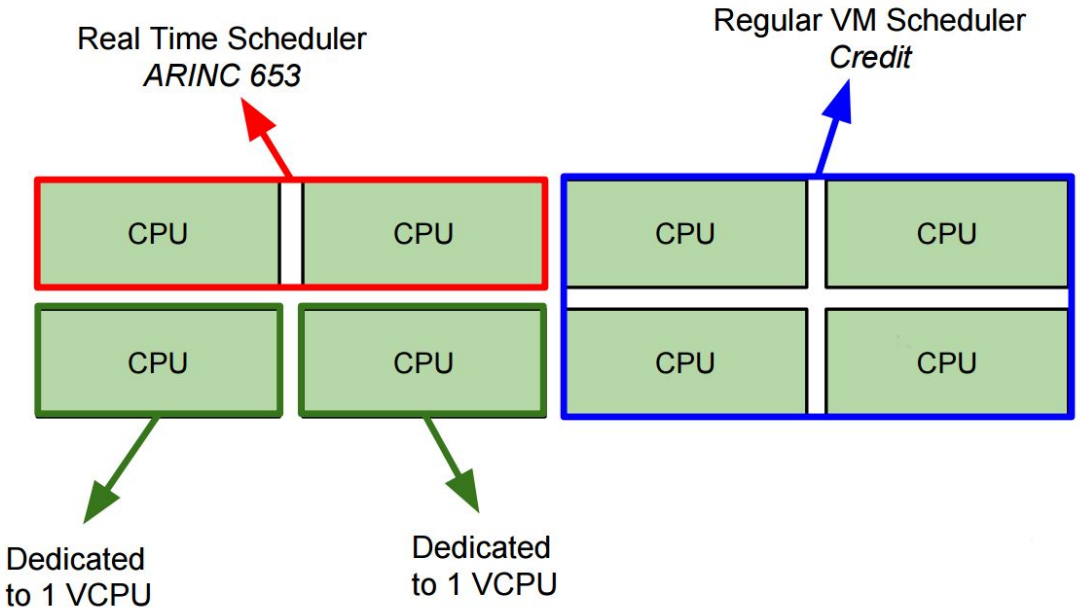
Real Time Scheduler
ARINC 653
Regular VM Scheduler
Credit
CPU
CPU
CPU
CPU
CPU
CPU
CPU
CPU
Dedicated
to 1 VCPU
Dedicated
to 1 VCPU

“Dom0less”

- Multiple domains can be created from Xen
  - Reduce boot time for boot VMs
- Described in the Device-Tree
- Arm only feature (so far)
- The VM cannot be rebooted
- PV and pass-through device supported

```
chosen {
    domU1 {
        compatible = "xen,domain";
        #address-cells = <0x2>;
        #size-cells = <0x1>;
        memory = <0 131072>;
        cpus = <2>;
        vpl011;

        module@0x4a000000 {
            compatible = "multiboot,kernel", "multiboot,module";
            reg = <0x0 0x4a000000 0xffffff>;
            bootargs = "console=ttyAMA0 init=/bin/sh";
        };

        module@0x4b000000 {
            compatible = "multiboot,ramdisk", "multiboot,module";
            reg = <0x0 0x4b000000 0xffffff>;
        };
    };

    domU2 {
        compatible = "xen,domain";
        #address-cells = <0x2>;
        #size-cells = <0x1>;
        memory = <0 65536>;
        cpus = <1>;

        module@0x4c000000 {
            compatible = "multiboot,kernel", "multiboot,module";
            reg = <0x0 0x4c000000 0xffffff>;
            bootargs = "console=ttyAMA0 init=/bin/sh";
        };

        module@0x4d000000 {
            compatible = "multiboot,ramdisk", "multiboot,module";
            reg = <0x0 0x4d000000 0xffffff>;
        };
    };
};
```

Hyperlaunch

- Ability to configure from a boot domain
- Security focused (Principles of Least Privilege)
- Building upon dom0less
- Currently under review for x86

Static partioning

- Pre-defined physical regions
- Limited to domains created by Xen (so far)
- Allow identity mapping
  - Useful on IOMMU-less platform

```
/ {
    chosen {
        #address-cells = <0x1>;
        #size-cells = <0x1>;
        ...
        domU1 {
            compatible = "xen,domain";
            cpus = <2>;
            memory = <0x0 0x80000>;
            xen,static-mem = <0x30000000 0x20000000>;
            ...
        };
    };
};
```

ARMv8-R

- Real-time profile
  - Cores provide hard real-time and safety guarantee
- The hypervisor is using an Memory Protection Unit (MPU)
- The VM can either use an MMU or MPU
- On-going development from AMD and Arm
  - Early boot is currently under review
- Can be tested on QEMU

Memory Protection Unit

- Only protecting the region
  - Limited number of regions (max 255)
- Flat mapping (virtual address = physical address)
- Xen needs to be build at a user provided address
  - Technically Position Independent Xen would work
- Memory needs to be statically partioned

Cache coloring

- Reduce interference
- Currently under review

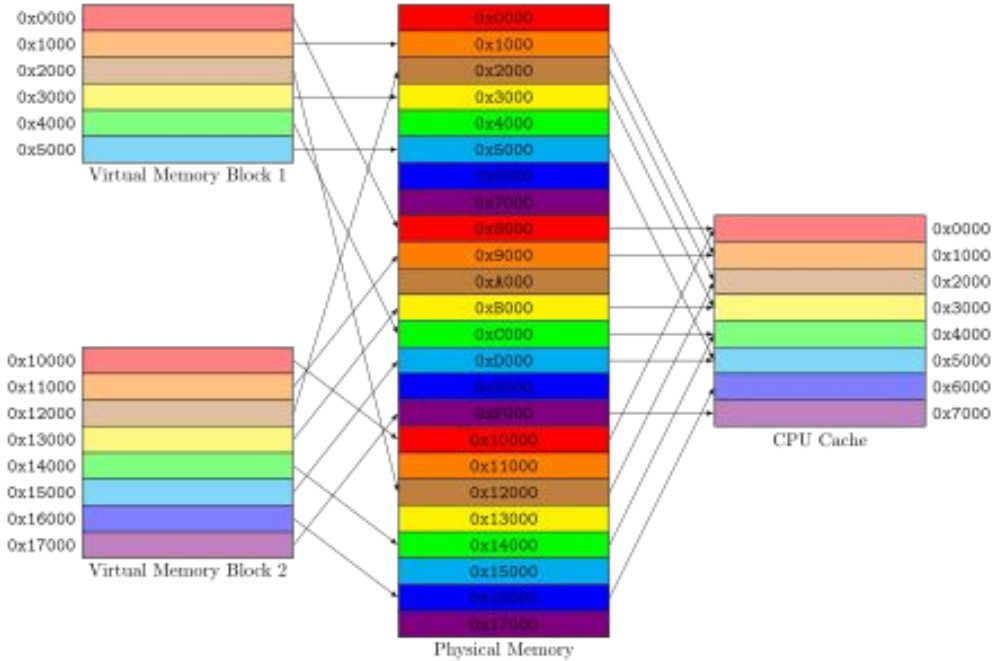

https://en.wikipedia.org/wiki/Cache_coloring

RTOS support

- FreeRTOS
- Zephyr
  - Guest
  - On-going development for Dom0

MISRA C

- Software guidelines
- Effort in Xen started in 2023
  - Help from Bugseng and sponsored by AMD
- Over 120 rules adopted
  - https://xenbits.xen.org/docs/unstable/misra/rules.html
- Use of ECLAIR to detect violations
  - Part of the gitlab CI

Requirements / Assumption of Use

- Part of safety certification
- Divided in 3 categories
  - Market requirements
    - High level functionalities
  - Product requirements
    - Specific concept and interfaces
  - Design requirements
    - Implementation details
- Requirements will be associated with tests

Xen Project Summit